Light Injection Operation and Data Rates

Philip Harris
1 November 2000

Basic assumptions

- The PMT gain is likely to be linear at low light levels, and will probably become non linear in the region of 100 photoelectrons. The exact shape of the gain curve therefore needs to be mapped out, and this is one of the basic functions of the light injection (LI) system. However, this is unlikely to change, at least on a short time scale. It is therefore planned that mapping the gain curve will take place no more than once per month. (In reality it may be much less often than this, but this needs to be established through experience).
- The gain itself may change over a time scale of hours, for example with daily temperature variations. It is therefore assumed to be necessary to monitor such changes on an hourly basis. However, as we are assuming that the shape of the gain curve is constant, only one point on that curve needs to be monitored to look for time variations of the gain.
- Although unlikely, it is not inconceivable that the gain variation with time may be different for different pixels. At the cost of factors of 3 (FD) and 6 (ND) in data rates, we should monitor all pixels for gain variations.
- It should not be necessary to monitor all WLS fibres within a pixel. However, in case it does become necessary, its impact is discussed in the addendum at the end of this note.

Modes of operation

a) Mapping out the gain curve involves injecting light at ten or so points along the curve. On a monthly basis or less, quantities of data required are relatively small.

b) Monitoring short-term gain changes involves flashing a number of LEDs, each about 1000 times at one light level up to once per hour. The actual LEDs used will be rotated, perhaps on a monthly basis, in order to avoid burning out particular spots on the M16s. In the near detector all LEDs have to be flashed in order to illuminate all pixels.

Far detector data rates

Let us first consider the extreme case:

- There are 20 LEDs illuminating each plane. With just three of these, however, the entire set of pixels may be illuminated. This gives a total of 48 LEDs that have to be flashed on an hourly-calibration basis, giving 70 seconds, i.e. just over one minute, per LED.

- An appropriate level at which to flash is probably in the range 30-70 photoelectrons. This matches the nonuniformity of the light injection system, while staying at reasonable light levels within the linear region.
- A level of 30 photoelectrons has a spread of approximately 20% from Poisson statistics. In principle 400 flashes would be adequate to establish the mean to 1%, but allowing for a safety margin suggests that we should have ~1000 flashes per data point.
- Each flash illuminates up to 640 WLS fibres, and therefore up to 1280 M16 pixels (counting both ends of each fibre).
- Cross talk could result in signals from all, instead of half, of the pixels in an M16 being illuminated. This gives a maximum of 2560 pixels on one flash.
- Each hit corresponds to 8 bytes of data. This gives up to 20 kB for each flash.
- In the trigger farm, all sets of LI events that arrive within a timeframe (~ 1 s) are analysed to give, e.g., mean and RMS of response of each channel, and summary events are written out. In order to be efficient, the number of events included in each such summary should be as large as possible.
- We should aim to stay well below the 2.5 Mb/s instantaneous rate limit of each crate: a reasonable limit to aim for is 1 Mb/s spread over the 2 crates, which limits us to 50 flashes in a 1 s timeframe.

Conclusion: in far detector, monitoring a single point on the gain curve each hour will take the form of 20 bunches of 50 flashes, with each set of 50 flashes to be within a second. The flasher will not put more than 1 Mb of data into a timeframe. If this limit allows, more than 50 flashes will be included in the timeframe to improve the efficiency of data reduction in the trigger farm analysis algorithm.

In this scenario, flashing would occupy 20 seconds of each minute. Mapping out the gain curve involves approximately 60 times as much data, and thus could be carried out for another 20 seconds of each minute over a period of 60 hours.

Near detector data rates

- Each LED illuminates a set of up to 10 strips in up to 60 planes in the calo section, and in all 33 planes in the spectrometer section. The data rate bottleneck occurs in the readout crates, of which there are 6 in the calo section and 2 in the spectrometer section. Thus, a spectrometer section readout crate sees 17 planes' worth of 10 strips per plane when an LED flashes, i.e. up to 170 strips.
- Assuming crosstalk, we may see signals in all M64 pixels instead of just the 10 targeted.
- From data taken at CERN over the summer, it is seen that signal pulses (both flasher and data) have tails extending up to about 50 ns, from the WLS fibre. The reflector connector will add a signal that is delayed by about 40 ns. Thus each hit may occupy 5 timeslices (although the crosstalk signals, being much smaller, are likely to fall below threshold more quickly and to occupy fewer timeslices).

- Each timeslice is 8 bytes long. Each flash can therefore occupy about 120 kB of memory, of which up to about 40 kB will pass through one readout crate. In order to stay at or around 1 Mb/s per crate, we should therefore flash at ~25 Hz (more if cross talk, pulse lengths etc allow), thereby completing the requisite 1000 flashes in ~40 seconds. With 33 LEDs to flash, this occupies less than 25 minutes within the one hour allowed. This would put up to 3 Mb of flasher data into a timeframe; much more than this would be difficult for the trigger farm processor memory to accommodate.

Conclusion for near detector: In this scenario, we would be flashing at ~25 Hz for just under half of the time. Mapping out the gain curve requires approximately 10 times as much data as a single point, as we will already be flashing all LEDs; this could therefore be done in about 10 hours by flashing more or less continuously.

Calibration detector

The calibration detector has just 6 LEDs, each illuminating up to 10 strips on 60 planes. The configuration is therefore similar to the near detector, but as there are fewer LEDs there would be much more time between bunches of flashes.

Reality check

In reality, rates will probably be reduced substantially. There are three likely mechanisms for this reduction:

- We will probably find that cross talk effects are less severe than described here, so the volume of data from each hit in the ND might well be a factor of 2 lower than estimated, allowing us to double the pulse rate and to halve the time spent pulsing.
- We will probably likewise find that it will not be necessary to pulse every pixel of each PMT, as the gain variation is likely to be by PMT rather than by pixel. This would reduce near-detector data rates by a factor 6, and far rates by a factor 2.
- We may well not have to monitor gain changes every hour; if the system is sufficiently stable, every few hours may be adequate.

Overall, therefore, rates may drop by an order of magnitude or more from the worst-case scenario described above. We cannot be certain though until the detector is in place, so we have to ensure that we have the ability to deal with the relatively high rates that may occur.

Deadtime

Far detector: roughly 5 microseconds per event; so even if we flash every LED 1,000 times every hour, that gives 32,000 flashes * 5E-6 sec/flash = 0.16 sec/hour.
Near detector: Suppose that, because of different length fibres etc., each event lasts 200 ns. There will be 33 LEDs flashing 1000 times per hour, giving 33,000 * 2E-7 = 0.006 seconds every hour.
In both cases, dead time is negligible.

Database

The impact upon the database is currently being assessed: Oracle can handle these data flows without a problem, but the database update system may need some redesigning if it is to handle calibration constants that change on an hourly basis.

Conclusion

A scenario has been outlined in which the gain variations of all PMT pixels can be tracked over time on an hourly basis, whilst staying within the data rate constraints of the trigger farm. Although for the near detector in particular these rates are uncomfortably close to the data-rate limit, it is likely that the actual rates needed will be substantially lower.

Addendum

It should not be necessary to monitor all WLS fibres that are multiplexed onto each pixel on an hourly basis, as the gain change should be a function of the pixel as a whole, if not of the entire PMT. If however we do decide to monitor every strip on an hourly basis in this way, the far detector data rates would increase by a factor of 20/3. This would have to be accommodated (a) by running for 40 s out of each minute at twice the above rate, i.e. at 40% of maximum instantaneous rate, and (b) reducing the precision with which we measure each point (600 instead of 1000 flashes). This would give a constant throughput of about 1 Mb/s in each crate, leaving little room for manoeuvre, but it is just about possible in principle.